

Answering $n^{2+o(1)}$ Counting Queries with Differential Privacy is Hard

Jonathan Ullman*

School of Engineering and Applied Sciences
Harvard University, Cambridge, MA
jullman@seas.harvard.edu

July 31, 2012

Abstract

A central problem in differentially private data analysis is how to design efficient algorithms capable of answering large numbers of *counting queries* on a sensitive database. Counting queries of the form “What fraction of individual records in the database satisfy the property q ?” We prove that if one-way functions exist, then there is no algorithm that takes as input a database $D \in (\{0, 1\}^d)^n$, and $k = \tilde{\Theta}(n^2)$ arbitrary efficiently computable counting queries, runs in time $\text{poly}(d, n)$, and returns an approximate answer to each query, while satisfying differential privacy. We also consider the complexity of answering “simple” counting queries, and make some progress in this direction by showing that the above result holds even when we require that the queries are computable by constant depth (AC^0) circuits.

Our result is almost tight in the sense that nearly n^2 counting queries can be answered efficiently while satisfying differential privacy. Moreover, super-polynomially many queries can be answered in exponential time.

We prove our results by extending the connection between differentially private counting query release and cryptographic traitor-tracing schemes to the setting where the queries are given to the sanitizer as input, and by constructing a traitor-tracing scheme that is secure in this setting.

1 Introduction

Consider a database $D \in (\{0, 1\}^d)^n$, in which each of the n rows corresponds to an individual’s record, and each record consists of d binary attributes. The goal of privacy-preserving data analysis is to enable rich statistical analyses on the database while protecting the privacy of the individuals. It is especially desirable to preserve *differential privacy* [DMNS06], which guarantees that no individual’s data has a significant influence on the information released about the database.

Some of the most basic statistics on a database are *counting queries*, which are queries of the form, “What fraction of individual records in D satisfy some property q ?” In particular we would like a differentially private *sanitizer* that, given a database D and k counting queries q_1, \dots, q_k from

*<http://seas.harvard.edu/~jullman>. Supported by NSF grant CNS-0831289 and a gift from Google, Inc.

the family \mathcal{Q} , outputs an approximate answer to each of the queries. We would like the number of queries, k , to be as large as possible, and the set of feasible queries, \mathcal{Q} , to be as general as possible. Ideally, \mathcal{Q} , would contain all counting queries.¹ Moreover, we would like the algorithm to run as efficiently as possible.

Early work in differential privacy [DN03, BDMN05, DMNS06] gave an efficient sanitizer—the so-called *Laplace Mechanism*. The Laplace Mechanism answers any set of k arbitrary efficiently computable counting queries by perturbing the answers with appropriately calibrated noise, proving good accuracy (say, within ± 0.01 of the true answer) as long as $k \lesssim n^2$.

The ability to approximately answer n^2 counting queries is indeed quite powerful, especially in settings where data is abundant and n is large. However, being limited to n^2 queries can be restrictive in settings where data is expensive or otherwise difficult to acquire, and n is small. It can also be restrictive when the budget of queries is shared between multiple analysts. Fortunately, a remarkable result of Blum et al. [BLR08] (with subsequent developments in [DNR⁺09, DRV10, HLM10]), showed that differentially private algorithms are not limited to n^2 queries. They showed how to approximately answer arbitrary counting queries even when k is *exponentially larger* than n . Unfortunately, their algorithm, and all subsequent algorithms capable of answering more than n^2 arbitrary counting queries, run in time (at least) $\text{poly}(2^d, n, k)$.

The result of Blum et al., raises the exciting possibility of an *efficient* algorithm that can privately compute approximate answers to large numbers of counting queries. Unfortunately, Dwork et al. [DNR⁺09] gave evidence that efficient sanitizers are inherently less powerful than their computationally unbounded counterparts. They consider the related problem of *counting query release*. In the counting query release problem, the goal is to produce a summary from which approximate answers to *every* query in \mathcal{Q} can be computed, and we'd like the sanitizer and the summary both to run in time much less than the size of \mathcal{Q} . In this setting, Dwork et al. constructed a family of roughly $2^{\sqrt{n}}$ queries that, under certain cryptographic assumptions, cannot be released efficiently (in time $\text{poly}(d, n)$), even though any family of size at most $\sim 2^n$ can be released by a computational unbounded algorithm. For any family \mathcal{Q} , efficiently solving the counting query release problem for \mathcal{Q} implies an efficient sanitizer for any polynomial number of queries from \mathcal{Q} . (See the related work for more discussion of this relationship.) Thus, hardness results for counting query release rule out a particular way of constructing efficient sanitizers. However, ultimately an analyst will only be able to ask a polynomial number of queries, and impossibility results for counting query release still leave room for optimism that there might be an efficient sanitizer that can answer many more arbitrary counting queries than the Laplace Mechanism.

Unfortunately, we show that this is not the case. We show that there is no efficient, differentially private algorithm that takes a database $D \in (\{0, 1\}^d)^n$, and $\tilde{\Theta}(n^2)$ arbitrary, efficiently computable counting queries as input and outputs an approximate answer to each of the queries. One way to summarize our results is that, unless we restrict the set of allowable queries \mathcal{Q} , or allow exponential running time, then the Laplace Mechanism is essentially the best possible algorithm for answering counting queries.

¹It may require super-polynomial time just to evaluate an arbitrary counting query, which would rule out efficiency for reasons that have nothing to do with privacy. For this discussion, we will always assume that the queries are efficiently computable, and are not the bottleneck in the computation.

1.1 Our Results and Techniques

In this paper we give new hardness results for answering counting queries while satisfying differential privacy. To make the statement of our results more concrete, we will assume that the counting queries are given to the sanitizer as input in the form of circuits that, on input an individual record $x \in \{0, 1\}^d$, decide whether or not the record x satisfies the property q . We say the queries are efficiently computable if the corresponding circuits are of size $\text{poly}(d, n)$.

Theorem 1.1. *Assuming the existence of one-way functions, there is no algorithm that, on input a database $D \in (\{0, 1\}^d)^n$ and $\tilde{\Theta}(n^2)$ efficiently computable counting queries, runs in time $\text{poly}(d, n)$ and returns an approximate answer to each query to within $\pm.49$, while satisfying differential privacy.*

We also show that, that the same theorem holds even for queries that are computable by unbounded-fan-in circuits of depth-6 over the basis $\{\wedge, \vee, \neg\}$ (a subset of the well-studied class AC^0), albeit under a strong (but still plausible) cryptographic assumptions.

Theorem 1.2. *Under the assumptions described in Section 5.2), there is no algorithm that, on input a database $D \in (\{0, 1\}^d)^n$ and $\tilde{\Theta}(n^2)$ efficiently computable depth-6 queries (circuits), runs in time $\text{poly}(d, n)$ and returns an approximate answer to each query to within $\pm.49$, while satisfying differential privacy.*

We now describe the techniques required to prove our results.

The Connection with Traitor-Tracing We prove our results by building on the connection between sanitizers for counting queries and *traitor-tracing schemes* utilized by Dwork et al. [DNR⁺09]. Traitor-tracing schemes were introduced by Chor, Fiat, and Naor [CFN94] for the purpose of identifying pirates who violate copyright restrictions. Roughly speaking, a (fully collusion-resilient) traitor-tracing scheme allows a sender to generate keys for n users so that 1) the sender can broadcast encrypted messages that can be decrypted by any user 2) any *efficient pirate decoder* capable of decrypting messages can be *traced* to at least one of the users who contributed a key to it, even if an arbitrary coalition of the users got together to contribute their keys.

Dwork et al. show that the existence of traitor-tracing schemes implies hardness results for the counting query release problem. Very informally, their argument is as follows: Suppose a coalition of users takes their keys and builds a database $D \in (\{0, 1\}^d)^n$ where each record contains one of their user keys. The family \mathcal{Q} will contain a query q_c for each possible ciphertext c . The query will ask “What fraction of the records (user keys) in D will decrypt the ciphertext c to the message 1?” Every user can decrypt, so if the sender encrypts a (single-bit) message m as a ciphertext c , then every user will decrypt c to m . Thus the answer to the counting query, q_c , will be m .

Suppose there were an efficient algorithm that could release the family \mathcal{Q} . Then the coalition could use it to efficiently produce a summary of the database D that enables them to efficiently compute an approximate answer to every query q_c , which would also allow them to efficiently decrypt the message. Such a summary will be an efficient pirate decoder, and thus the tracing algorithm can use its answers to identify one of the users in the coalition. However, if there is a way to identify one of the users in the database from the summary, then the summary cannot be differentially private.

In order to instantiate their result, they need a traitor-tracing scheme. Since \mathcal{Q} contains a query for every ciphertext, the parameter to optimize is the length of the ciphertexts. Using the

fully collusion-resilient traitor-tracing scheme of Boneh, Sahai, and Waters [BSW06], which has ciphertexts of length $\sim \sqrt{n}$, they obtain a family of queries of size $\sim 2^{\sqrt{n}}$ that cannot be released efficiently.

Our Approach In our setting, we don't expect to answer every query in \mathcal{Q} , only the set of $k \ll |\mathcal{Q}|$ input queries. At first glance, this should make answering the queries much easier, and thus make it more difficult to demonstrate hardness. However, the attacker does have the power to choose the queries he wants answers to, and can choose just the queries that are most harmful to privacy. Our first observation is that in the traitor-tracing scenario, the tracing algorithms will only query the pirate decoder on a polynomial number of ciphertexts, which will be randomly chosen and depend on the particular keys that were instantiated for the scheme. For many schemes, even $\tilde{O}(n^2)$ queries is sufficient. Thus it would seem that the tracing algorithm could simply decide which queries it will make, give those queries as input to the sanitizer, and then use the answers to those queries to identify a user and violate differential privacy.

Although this observation would seem to be sufficient to establish our main result, the intuition we sketched ignored an important issue. Many traitor-tracing schemes (including that of [BSW06]) are only able to trace *stateless* pirate decoders, which essentially commit to a response to each possible query (or a distribution over responses) once and for all. For the counting query release problem, the private summary is necessarily stateless, and thus the result of Dwork et al. can be instantiated with any scheme that allows tracing of stateless pirate decoders. However, the type of sanitizer might give answers that depend on the whole sequence of queries, and is given all the queries it will have to answer at once. Thus, in order to prove our results, we will need a traitor-tracing scheme that can trace *stateful* pirate decoders, and selects all its queries to the pirate decoder at once.

The problem of tracing stateful pirates is quite natural even without the implications for private data analysis. To be sure, this problem has been studied in the literature, originally by Kiayias and Yung [KY01]. However, their solution, and all others known, does not apply to our specific setting (see discussion below). However, we also refine the basic connection between traitor-tracing schemes and differential privacy by showing that, in many respects, fairly weak traitor-tracing schemes suffice to establish the hardness of preserving privacy. In particular, although the pirate decoder may be stateful, it will be constrained in other ways, which will make it easier to construct a traitor-tracing scheme for these kinds of pirates. Indeed, we construct such a scheme to establish Theorem 1.1. The scheme will also have weakened requirements in other respects, having nothing to do with the statefulness of the pirate. These weakened requirements allow us to reduce the complexity of the decryption, which means that the queries used by the attacker do not need to be arbitrary polynomial-size circuits, but instead can be circuits of constant-depth, which will establish Theorem 1.2. See Sections 3.1 and 4 for a precise statement of the kind of traitor-tracing scheme that will suffice and Section 5 for our construction.

1.2 Related Work

Traitor-Tracing Schemes Chor, Fiat, and Naor [CFN94] introduced the notion of a traitor-tracing scheme, and they have been studied extensively as a means of distributing copyrighted content. The connection between traitor-tracing schemes and hardness results for differentially private sanitizers (discussed above) was discovered by Dwork et al. [DNR⁺09]. The literature

on traitor-tracing schemes is too vast to summarize here, and much of it focuses on constructing schemes with properties that are not relevant to our application.

The work that is most related to ours is that of Kiayias and Yung [KY01], which considers traitor-tracing schemes that can trace stateful pirate decoders. However, their construction relies on a certain “watermarking assumption” that makes sense in the context of distributing copyrighted content, but does not apply when the messages are single bits, as is the case in our application. However, our scheme relies on stronger assumptions about the pirate decoder (that only make sense in the context of hardness results for differential privacy) than their does, making our results incomparable to theirs.

Fingerprinting codes—an ingredient of our traitor-tracing scheme—were introduced by Boneh and Shaw [BS98], also for the problem of watermarking copyrighted content. Fingerprinting codes have been used extensively in constructions of traitor-tracing schemes (cf. Boneh and Naor [BN08]). To achieve the best parameters for our scheme, we use a construction of fingerprinting codes of optimal length, due to Tardos [Tar08].

The Relationship with [DNV12] Dwork, Naor, and Vadhan [DNV12] gave information theoretic lower bounds for *stateless sanitizers*. These are sanitizers that take k queries as input, but whose answers to each query do not depend on the other $k - 1$ input queries. They showed that (even computationally unbounded) stateless sanitizers can answer at most $\sim n^2$ queries with non-trivial accuracy, while satisfying differential privacy. The Laplace Mechanism is a stateless sanitizer that answers $\sim n^2$ queries, and thus their result is tight in this respect.

Another interpretation of our results, that would lead to an alternative proof of our results, is that we construct a family of queries for which “keeping state doesn’t help”. Unfortunately, the use of traitor-tracing schemes and fingerprinting codes in our proofs obscures the intuitive relationship between our arguments and those in [DNV12], so we will give an informal discussion here.

They consider a game where an n -row database is chosen at random, and a random subset of $n - 1$ of those rows is given to the attacker. The attacker wants to violate privacy by recovering the n -th row. To do so, the attacker chooses $\sim n^2$ queries (randomly, from a distribution that depends on the $n - 1$ known rows) and requests answers to these queries. Using these answers, they show that there is a particular way for the attacker to (randomly) guess the missing row, that will succeed with sufficient probability to constitute a privacy violation. Their argument is in two steps: 1) The expected correlation between the answers given by a stateless sanitizer and the answers on a database consisting only of the missing row is significant. 2) A stateless sanitizer cannot give answers that are correlated with too many rows that are not in the database. Combining these two steps shows that the attacker has a significant chance of identifying the n -th row from its correlation with the answers.

Typically, the intuition behind the analysis of traitor-tracing schemes follows roughly the same two steps: 1) There will be some correlation between the decryptions returned by the efficient pirate and the decryptions that would be returned by some member of the coalition (using only his own key). 2) There will not be significant correlation between the decryptions returned by the efficient pirate and the decryptions that would be returned by any user not a member of the coalition. Indeed, if we “unrolled” the analysis of the fingerprinting code directly into our construction, we would make exactly the same arguments.

Other Types of Sanitizers There are two other variants of the counting query problem that have appeared in the literature. The first, which we have already discussed, is counting query release. The second, is *interactive sanitization*. This problem is the same as the one we consider, where the sanitizer is given a database D and k queries from a family \mathcal{Q} , however the queries arrive one at a time, and may be chosen adaptively. In this setting, we want the sanitizer to answer each query efficiently (in time polynomial in d , n , and k). The Laplace Mechanism is, in fact, interactive, and a line of work initiated by Roth and Roughgarden as well as Hardt and Rothblum [RR10, HR10, GRU12] showed how to interactively answer nearly 2^n queries in time $\text{poly}(2^d, n)$ per query.

The three variants we’ve described satisfy some interesting relationships. First, if we have an algorithm that runs in time T and releases a summary that enables an analyst to answer any query in \mathcal{Q} in time T , then we also have an interactive sanitizer that runs in time T per query that answers any k queries from \mathcal{Q} . Secondly, if we have an interactive sanitizer that answers k queries from \mathcal{Q} in time T per query, then we also have a non-interactive sanitizer that answers K queries from \mathcal{Q} in time Tk . Thus, holding \mathcal{Q} fixed (and assuming $k = \text{poly}(d, n)$), the problem we consider is the easiest form of private counting query release, and the lower bounds we prove imply lower bounds for the other variants.

For the case of interactive sanitization, these lower bounds are new. To our knowledge, prior to our work it was possible that there was an interactive sanitizer that ran in time $\text{poly}(d, n)$ per query and answered nearly 2^n arbitrary counting queries, whereas our results show that there is no efficient interactive sanitizer for significantly more than n^2 queries. On the other hand, for counting query release, our results only show that it is hard to release a particular family of queries \mathcal{Q} whose size is at least 2^n . For families of queries this large, the results of Dinur and Nissim [DN03] already imply the impossibility of release, even by computationally unbounded algorithms.

We note that typically \mathcal{Q} is not the same in the data release problem as it is for sanitizers. For data release, \mathcal{Q} must be smaller than 2^n , and thus cannot contain all efficiently computable counting queries. Sanitizers (both interactive and non-interactive) are supposed to circumvent this problem by allowing the queries to be arbitrary, but only answering the k queries that are needed. However our results show that they can only circumvent the problem if $k \ll n^2$ queries is sufficient.

Hardness Results for Synthetic Data There has been considerable focus in differentially private data analysis on sanitizers that produce *synthetic data* [BCD⁺07, BLR08, DNR⁺09, DRV10]. A sanitizer outputs synthetic data if on input a database $D \in (\{0, 1\}^d)^n$, it outputs a new database $\hat{D} \in (\{0, 1\}^d)^{\hat{n}}$ that approximately preserves the answer to each of some set of queries. In addition to being a natural and well-studied desideratum for private data analysis, essentially all known techniques for answering $\gg n^2$ queries (from sufficiently general families \mathcal{Q}) output synthetic data. Even many constructions of interactive sanitizers for answering $\gg n^2$ queries [RR10, HR10, GRU12], can easily be modified to output a synthetic database. However, even the best of these mechanisms run in time $\text{poly}(2^d, n, k)$.

Unfortunately, Ullman and Vadhan [UV11], building on work by Dwork et al. [DNR⁺09] showed that exponential running time is inherent for sanitizers that output synthetic data, even if the synthetic database only has to preserve the answers to *2-way marginals* (roughly, the mean of each column and the pairwise correlation between columns). These results apply even when the number of queries is $\ll n^2$, and thus apply to problems where efficient algorithms that do not output synthetic data (e.g. the Laplace mechanism) are known. Thus, these results say more about

the hardness of generating synthetic data, and the limitations of current techniques, than they do about the hardness of answering large numbers of counting queries.

Answering Simple Counting Queries There has been a significant body of research on designing improved algorithms for releasing “simple” families of queries (which, as discussed above, implies interactive and non-interactive sanitizers for these families of queries). For instance, Hardt, Rothblum, and Servedio [HRS12] as well as Thaler, Ullman, and Vadhan [TUV12] recently gave algorithms for releasing the family of monotone k -way conjunctions. A monotone k -way conjunction is specified by a subset of the columns, $S \subseteq [d]$, $|S| = k$, and asks “What fraction of records in D have a 1 in every column in S ?” Note that there are $\sim d^k$ such queries (for $k \ll d$). These queries have been identified as an especially important family for differentially private data release (cf. [BCD+07, KRSU10, GHRU11]) The two works mentioned give efficient algorithms capable of releasing all monotone k -way conjunctions on a database of size $d^{O(\sqrt{k})}$, and thus are capable of answering $n^{\Omega(\sqrt{k})} \gg n^2$ queries from this family (for a particular choice of n).

Thus there is a significant gap between k -way conjunctions, for which there are efficient, non-trivial improvements on the Laplace Mechanism, and polynomial-size depth-6 circuits, for which we can show there is no efficient algorithm that significantly improves on the Laplace Mechanism.

2 Preliminaries

Differentially Private Algorithms Let a *database* $D \in (\{0, 1\}^d)^n$ be a collection of n rows $(x^{(1)}, \dots, x^{(n)}) \in \{0, 1\}^d$. We say that two databases $D, D' \in (\{0, 1\}^d)^n$ are *adjacent* if they differ only on a single row, and we denote this by $D \sim D'$. Let $\mathcal{M}: (\{0, 1\}^d)^n \rightarrow \mathcal{R}$ be a randomized algorithm that takes a database as input. For ease of notation, we will write \mathcal{M} as a function on databases in $(\{0, 1\}^d)^n$, but we will always think of the sanitizer as taking databases of arbitrary dimensions as input, and thus the parameters in various definitions may depend on the dimensions of the input.

Definition 2.1 (Differential Privacy [DMNS06]). An algorithm \mathcal{M} is (ε, δ) -*differentially private* if for every two adjacent databases $D \sim D' \in (\{0, 1\}^d)^n$ and every subset $S \subseteq \mathcal{R}$,

$$\Pr[\mathcal{M}(D) \in S] \leq e^\varepsilon \Pr[\mathcal{M}(D') \in S] + \delta.$$

If \mathcal{M} is (ε, δ) -differentially private for some functions $\varepsilon = \varepsilon(n) = O(1)$, $\delta = \delta(n) = o(1/n)$, we will drop the parameters ε and δ and say that \mathcal{M} is *differentially private*.

The choice of $\varepsilon = O(1)$, $\delta = o(1/n)$ is a fairly weak setting of the privacy parameters, and most known constructions of differentially private mechanisms for various tasks give quantitatively stronger privacy guarantees. These parameters are essentially the weakest parameters possible, as (ε, δ) -differentially privacy is not a meaningful privacy guarantee for $\varepsilon = \omega(1)$ or $\delta = \Omega(1/n)$. The fact that our lower bounds apply to these parameters makes our results stronger.

Sanitizers for Counting Queries Since an algorithm that always outputs \perp satisfies Definition 2.1, we also need to specify what it means for the sanitizer to be useful. In this paper we study sanitizers that give accurate answers to *counting queries*. A counting query on $\{0, 1\}^d$ is defined

by a predicate $q: \{0, 1\}^d \rightarrow \{0, 1\}$. Abusing notation, we define the evaluation of the query q on a database $D \in (\{0, 1\}^d)^n$ to be

$$q(D) = \frac{1}{n} \sum_{i=1}^n q(x^{(i)})$$

We will use $\mathcal{Q}^{(d)}$ to denote a set of counting queries on $\{0, 1\}^d$ and $\mathcal{Q} = \bigcup_{d \in \mathbb{N}} \mathcal{Q}^{(d)}$.

We are interested in sanitizers that take a sequence of queries from some set \mathcal{Q} as input, which we call *generic sanitizers*. Formally a generic sanitizer is a function $\mathcal{M}: (\{0, 1\}^d)^n \times (\mathcal{Q}^{(d)})^k \rightarrow \mathbb{R}^k$. Notice that we assume that \mathcal{M} outputs k real-valued answers. Think of the j -th component of the output of \mathcal{M} as an answer to the j -th query. For the results in this paper, this assumption will be without loss of generality.² Again, for ease of notation, we will write \mathcal{M} as a function on k queries in $\mathcal{Q}^{(d)}$, but we will always think of the sanitizer as taking a database of arbitrary dimensions $n \times d$ and an arbitrary number of queries from $\mathcal{Q}^{(d)}$ as input. Thus the parameters in various definitions may depend on the dimensions of the input and the number of queries. Definition 2.1 extends naturally to generic sanitizers by requiring that for every $q_1, \dots, q_k \in \mathcal{Q}$, the sanitizer $\mathcal{M}_{q_1, \dots, q_k}(\cdot) = \mathcal{M}(\cdot, q_1, \dots, q_k)$ is (ϵ, δ) -differentially private as a function of the input database.

Now we formally define what it means for a generic sanitizer to give accurate answers.

Definition 2.2 (Accuracy). Let D be a database and q_1, \dots, q_k be a set of counting queries. A sequence of answers a_1, \dots, a_k is α -accurate for D and q_1, \dots, q_k if

$$\max_{j \in [k]} |q_j(D) - a_j| \leq \alpha.$$

Let \mathcal{Q} be a set of counting queries, $k \in \mathbb{N}$ and $\alpha, \beta \in [0, 1]$ be parameters. A generic sanitizer \mathcal{M} is $(\alpha, \beta, \mathcal{Q}, k)$ -accurate if for every database $D \in (\{0, 1\}^d)^n$ and every sequence of queries $q_1, \dots, q_k \in \mathcal{Q}^{(d)}$

$$\Pr_{\mathcal{M}'\text{'s coins}} [\mathcal{M}(D, q_1, \dots, q_k) \text{ is } \alpha\text{-accurate for } D \text{ and } q_1, \dots, q_k] \geq 1 - \beta.$$

If \mathcal{M} is $(\alpha, \beta, \mathcal{Q}, k)$ -accurate for any (constant) $\alpha < 1/2$ and $\beta = \beta(n) = o(1/n^2)$, we will drop α and β and say that \mathcal{M} is (\mathcal{Q}, k) -accurate.

The choice of $\alpha < 1/2, \beta = o(1/n^2)$ is also significantly weaker than what can be achieved by known constructions of generic sanitizers. These parameters are also essentially the weakest parameters possible, as a mechanism that answers $1/2$ to every query achieves $\alpha = 1/2, \beta = 0$ for any number of arbitrary queries. Also, if there is a mechanism that achieves $(< 1/2, \beta, \mathcal{Q}, k)$ -accuracy for $\beta > 1/2$, then there is another mechanism that achieves $(< 1/2, o(1/n^2), \mathcal{Q}, k)$ -accuracy with only a small loss in the privacy parameters and the efficiency of the mechanism. The fact that our lower bound applies to these parameters makes our results stronger.

² A certain settings, $\mathcal{M}(D, q_1, \dots, q_k)$ is allowed to output a “summary” $z \in \mathcal{R}$ for some range \mathcal{R} . In this case we would also require that there exists an “evaluator” $\mathcal{E}: \mathcal{R} \times \mathcal{Q} \rightarrow \mathbb{R}$ that takes a summary and a query and returns an answer $\mathcal{E}(z, q) = a$ that approximates $q(D)$. The extra generality is used to allow \mathcal{M} to run in less time than the number of queries its answering (e.g. releasing a fixed family of queries), but we can ignore this issue for our range of parameters. A generic sanitizer, \mathcal{M} that outputs a summary in \mathcal{R} can be converted to a generic sanitizer with output in \mathbb{R}^k simply by running $\mathcal{M}(D, q_1, \dots, q_k)$ to obtain $z \in \mathcal{R}$ and then computing $a_1 = \mathcal{E}(z, q_1), \dots, a_k = \mathcal{E}(z, q_k)$. This conversion increases the running time by a factor of k , which is the minimum time required to read the input queries. Thus, our assumption is without loss of generality.

Efficiency of Generic Sanitizers Simply, a generic sanitizer is efficient if it runs in time polynomial in the length of its input. To make the statement more precise, we need to specify how the queries are given to the sanitizer as input.

Notice that to specify an arbitrary counting query $q: \{0, 1\}^d \rightarrow \{0, 1\}$ requires 2^d bits and thus it may require time 2^d to evaluate. In this case, a sanitizer whose running time is polynomial in the time required to evaluate the query is not especially efficient; that is, its running time is not $\text{poly}(d, n, k)$. Thus, we restrict attention to queries that are efficiently computable, and have a succinct representation. In this work, we will fix the representation to be Boolean circuits over the basis $\{\wedge, \vee, \neg\}$ with possibly unbounded-fan-in. In this representation, a query whose description length is s can also be evaluated in time $\text{poly}(s)$.

For a function $s: \mathbb{N} \rightarrow \mathbb{N}$, we use $\mathcal{Q}_s^{(d)}$ to denote the set of all counting queries on $\{0, 1\}^d$ specified by circuits of size $s(d)$. In this work we also consider the case where the queries are computable by circuits of low depth. Similarly, for functions $s, h: \mathbb{N} \rightarrow \mathbb{N}$, we use $\mathcal{Q}_{s,h}^{(d)}$ to denote the set of all counting queries on $\{0, 1\}^d$ specified by circuits of size $s(d)$ and depth $h(d)$. Finally, we use $\mathcal{Q}_{\text{all}}^{(d)}$ to denote the set of all counting queries on $\{0, 1\}^d$.

Definition 2.3 (Efficient Generic Sanitizers). A generic sanitizer \mathcal{M} is *efficient* if, on input a database $D \in (\{0, 1\}^d)^n$ and k queries $q_1, \dots, q_k \in \mathcal{Q}_s^{(d)}$, \mathcal{M} runs in time $\text{poly}(d, n, k, s(d))$. A generic sanitizer \mathcal{M} is *efficient for depth- h queries* if, on input a database $D \in (\{0, 1\}^d)^n$ and k queries $q_1, \dots, q_k \in \mathcal{Q}_{s,h}^{(d)}$, \mathcal{M} runs in time $\text{poly}(d, n, k, s(d))$.

For comparison with our results, we will recall the properties of some known mechanisms:

Theorem 2.4 (Laplace Mechanism [DN03, DMNS06]). *There exists a generic sanitizer \mathcal{M}_{Lap} that is 1) differentially private, 2) efficient, and 3) $(\mathcal{Q}_{\text{all}}^{(d)}, \tilde{\Omega}(n^2))$ -accurate.*

Note that since this mechanism is efficient, it is also efficient for depth- h queries.

Theorem 2.5 (“Advanced Query Release Mechanisms” [BLR08, DNR⁺09, DRV10, HR10, HLM10]). *There exists a generic sanitizer \mathcal{M}_{Adv} that is 1) differentially private and 2) $(\mathcal{Q}_{\text{all}}^{(d)}, 2^{\tilde{\Omega}(n/\sqrt{d})})$ -accurate. For queries in $\mathcal{Q}_s^{(d)}$, \mathcal{M}_{Adv} runs in time $\text{poly}(2^d, n, k, s(d))$.*

As we mentioned above, these mechanisms can achieve stronger quantitative privacy and accuracy guarantees (in terms of ε, δ for privacy and α, β for accuracy) with only a small degradation in the number of queries. However, we state the results for our relaxed choice of parameters both for simplicity and for comparison with our negative results. Also, notice that both of these mechanisms provide accuracy guarantees that are independent of the complexity of the queries (although the running time of the mechanism will depend on the complexity of the queries). Our hardness results will apply to sanitizers that only provide accuracy for queries of size $\text{poly}(d, n)$.

3 Traitor-Tracing Schemes and Other Cryptographic Primitives

In this section we give a definitions of traitor-tracing schemes and of the other cryptographic primitives that will be useful in proving our result. For clarity, we will use \mathbf{A}_{TT} , and \mathbf{A}_{FP} to denote algorithms associated with traitor-tracing schemes and fingerprinting codes, respectively. Algorithms associated with encryption schemes will typically have no subscript, however we will use a_{Enc} for their associated parameters.

3.1 Traitor-Tracing Schemes

We now give a definition of a traitor-tracing scheme, heavily tailored to the task of proving hardness results for generic sanitizers. We will sacrifice some consistency with the standard definitions. See below for further discussion of the ways in which our definition departs from the standard definition of traitor-tracing. In some cases, the non-standard aspects of the definition will be necessary to establish our results, and in others it will be for convenience. Despite these differences, we will henceforth refer to schemes satisfying our definition simply as *traitor-tracing schemes*.

Intuitively, a traitor-tracing scheme is a form of broadcast encryption, in which a sender can broadcast an encrypted message that can be decrypted by each of a large set of users. The standard notion of security for such a scheme would require that an adversary that doesn't have any of the keys cannot decrypt the message. A traitor-tracing scheme has the additional property that given any decoder capable of decrypting the message (which must "know" some of the keys), there is a procedure for determining which user's key is being used. Moreover, we want the scheme to be "collusion resilient", in that even if a coalition of users gets together and combines their keys in some way to produce a decoder, there is still a procedure that identifies at least one member of the coalition.

We now describe the syntax of a traitor-tracing scheme more formally. For functions $n, k_{\text{TT}}: \mathbb{N} \rightarrow \mathbb{N}$, an (n, k_{TT}) -traitor-tracing scheme is a tuple of algorithms $(\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$. We allow all the algorithms to be randomized except for Dec_{TT} .³

- The algorithm Gen_{TT} takes a security parameter, κ , and returns a set of $n = n(\kappa)$ user keys $\vec{sk} = (sk^{(1)}, \dots, sk^{(n)}) \in \{0, 1\}^\kappa$. Formally, $\vec{sk} = (sk^{(1)}, \dots, sk^{(n)}) \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa)$.
- The algorithm Enc_{TT} takes a set of user keys \vec{sk} and a message bit $b \in \{0, 1\}$, and generates a ciphertext $c \in \mathcal{C} = \mathcal{C}^{(\kappa)}$. Formally, $c \leftarrow_{\text{R}} \text{Enc}_{\text{TT}}(\vec{sk}, b)$.
- The algorithm Dec_{TT} takes any single user key sk and a ciphertext $c \in \mathcal{C}$, runs in time $\text{poly}(\kappa, n(\kappa))$ and deterministically returns a message bit $\bar{b} \in \{0, 1\}$. Formally $\bar{b} = \text{Dec}_{\text{TT}}(sk, c)$.
- The algorithm trace takes a set of user keys $\vec{sk} \in (\{0, 1\}^\kappa)^{n(\kappa)}$ as regular input and an oracle $\mathcal{P}: (\mathcal{C}^{(\kappa)})^{k_{\text{TT}}(\kappa)} \rightarrow \{0, 1\}^{k_{\text{TT}}(\kappa)}$, makes at most $k_{\text{TT}} = k_{\text{TT}}(\kappa)$ non-adaptive queries $c_1, \dots, c_{k_{\text{TT}}} \in \mathcal{C}^{(\kappa)}$ to its oracle, and returns the name of a user $i \in [n(\kappa)]$. Formally, $i \leftarrow_{\text{R}} \text{Trace}_{\text{TT}}^{\mathcal{P}}(\vec{sk})$.

Intuitively, think of the oracle \mathcal{P} as being given some subset of keys $\vec{sk}_S = (sk^{(i)})_{i \in S}$ for a non-empty set $S \subseteq [n]$, and Trace_{TT} is attempting to identify a user $i \in S$. Clearly, if \mathcal{P} ignores its input and always returns 0, Trace_{TT} cannot have any hope of success, so we need to place some condition on \mathcal{P} that allows Trace_{TT} to succeed. Roughly, we want to require that Trace_{TT} is successfully decrypting messages encrypted under the keys \vec{sk} , however for convenience, we will place a stronger requirement on \mathcal{P} . Note that making stronger assumptions about \mathcal{P} can only help the tracing algorithm, so as long as the assumptions are still valid in the intended application, they cannot hurt. Specifically, we will place the following requirement on \mathcal{P} .

Definition 3.1 (Available Pirate Decoder). Let $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ be an (n, k_{TT}) -traitor-tracing scheme. Let \mathcal{P} be a (possibly randomized) algorithm. We say that \mathcal{P} is

³It would not substantially affect our results if Dec_{TT} were randomized, however we will assume that Dec_{TT} is deterministic for ease of presentation.

a k_{TT} -available pirate decoder if for every $\kappa \in \mathbb{N}$, every set of user keys $\vec{sk} = (sk^{(1)}, \dots, sk^{(n)}) \in \{0, 1\}^\kappa$, every $S \subseteq [n]$ such that $|S| \geq n - 1$, and every $c_1, \dots, c_{k_{\text{TT}}} \in \mathcal{C}^{(\kappa)}$,

$$\Pr \left[\begin{array}{c} (\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}}) \leftarrow_{\text{R}} \mathcal{P}(\vec{sk}_S, c_1, \dots, c_{k_{\text{TT}}}) \\ \exists j \in [k_{\text{TT}}], b \in \{0, 1\} \left((\forall i \in S, \text{Dec}_{\text{TT}}(sk^{(i)}, c_j) = b) \wedge (\bar{b}_j \neq b) \right) \end{array} \right] \leq o\left(\frac{1}{n(\kappa)^2}\right).$$

In other words, if every user key $sk^{(i)}$ (for $i \in S$) decrypts c to 1 (resp. 0), then $\mathcal{P}(\vec{sk}_S, \cdot)$ decrypts c to 1 (resp. 0).

We can now define a secure, (n, k_{TT}) -traitor-tracing scheme:

Definition 3.2 (Traitor-Tracing Scheme). Let $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ be an (n, k_{TT}) -traitor-tracing scheme. Let $k_{\text{TT}}: \mathbb{N} \rightarrow \mathbb{N}$ be a function. We say that Π_{TT} is a *secure* (n, k_{TT}) -traitor-tracing scheme if for every (possibly randomized) algorithm \mathcal{P} that 1) runs in $\text{poly}(\kappa, n(\kappa), k_{\text{TT}}(\kappa))$ time and 2) is a k_{TT} -available pirate decoder, and for every $S \subseteq [n(\kappa)]$ such that $|S| \geq n(\kappa) - 1$,

$$\Pr_{\substack{\vec{sk} \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa) \\ \mathcal{P}'\text{'s, Trace}_{\text{TT}}\text{'s coins}}} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}(\vec{sk}_S, \cdot)}(\vec{sk}) \notin S \right] = o\left(\frac{1}{100n(\kappa)}\right)$$

Remarks About Our Definition of Traitor-Tracing The traitor-tracing schemes we consider are somewhat different than those previously studied in the literature. Specifically:

- Our traitor-tracing schemes are *private key* in every respect. That is, we do not require the encryption or tracing algorithms to use public keys. In the typical application of traitor-tracing schemes to content distribution, these would be desirable features, however they are not necessary for this application. We take advantage of this relaxation in two ways: 1) Since we do not require a public-key cryptosystem, our first result (Theorem 1.1) only needs to assume the existence of one-way functions. 2) Since private-key cryptosystems are easier to construct, we are able to find a candidate scheme whose decryption can be implemented by constant-depth circuits, which we use to instantiate Theorem 1.2.
- We only require that the tracing algorithm succeeds with probability $1 - o(1/n) = 1 - 1/\text{poly}(\kappa)$ (in the most common regime where $n = n(\kappa) = \text{poly}(\kappa)$). Typical traitor-tracing schemes would require that the tracing algorithm succeeds with probability $1 - \text{negl}(\kappa)$. As above, we use this extra flexibility to find a candidate scheme whose decryption can be implemented by constant-depth circuits.
- We do not give the pirate decoder access to an encryption oracle. In other words, we do not require CPA security. Most traitor-tracing schemes in the literature are public-key, making this distinction irrelevant. Here, we only need an encryption scheme that is secure for an a priori bounded number of messages. As above, we use this extra flexibility to find a candidate scheme whose decryption can be implemented by constant-depth circuits.
- We do not require that the key generation, encryption, or tracing algorithms to be efficient, as the standard definition of differential privacy requires the privacy guarantees to hold for every database (which, roughly, will be generated by the key generation algorithm) and every set of queries (which, roughly, will correspond to the encryption and tracing algorithm). However,

we do require decryption to be efficient, as the queries used in our results must compute decryption, and we require that the queries are efficiently computable. In our constructions all of these algorithms will be efficient, but we are hopeful that the extra flexibility may be useful for proving stronger lower bounds for differentially private data analysis.

- We allow the pirate decoder to be *stateful*, but in an unusual way. The pirate is actually allowed to see all the queries at once, whereas typically even traitor-tracing schemes for stateless adversaries are allowed to query the pirate adaptively, and in this sense our model of a stateful pirate is more powerful than what has been previously considered. However, we do require (roughly) that if any of the queries are ciphertexts generated by $\text{Enc}(\vec{s}k, b)$, then the pirate decoder answers b to those queries, regardless of the other queries issued. In typical applications of traitor-tracing, the pirate could simply answer \perp to every query if it detects that any of them are malformed. Kiayias and Yung [KY01] addressed this problem in the context of content distribution, and showed a generic transformation from a traitor-tracing scheme that traces stateless pirates to one that traces stateful pirates (for the more standard notion of stateful pirates). Their approach relies on a particular “watermarking assumption” that cannot apply to bit-encryption, making their results incomparable to ours.

3.2 Fingerprinting Codes

To construct a traitor-tracing scheme that satisfies Definition 3.2, we will make use of (*binary*) *fingerprinting codes*, introduced by Boneh and Shaw [BS98]. A fingerprinting code is a pair of efficient (possibly randomized) algorithms $(\text{Gen}_{\text{FP}}, \text{Trace}_{\text{FP}})$ with the following syntax.

- The algorithm Gen_{FP} takes a number of users n as input and outputs a codebook of n codewords of length $\ell_{\text{FP}} = \ell_{\text{FP}}(n)$, $W = (w^{(1)}, \dots, w^{(n)}) \in \{0, 1\}^{\ell_{\text{FP}}}$. Formally $W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n)$. We will typically treat $W \in \{0, 1\}^{n \times \ell_{\text{FP}}}$ as a matrix with each row containing a codeword.
- The algorithm Trace_{FP} takes an n -user codebook W and a word $w' \in \{0, 1\}^{\ell_{\text{FP}}}$ and returns an index $i \in [n]$. Formally, $i = \text{Trace}_{\text{FP}}(W, w')$.

Given a non-empty subset $S \subseteq [n]$ and a set of codewords $W_S = (w^{(i)})_{i \in S} \in \{0, 1\}^{\ell_{\text{FP}}}$, we define the set of *feasible codewords* to be

$$F(W_S) = \left\{ w' \mid \forall j \in [\ell_{\text{FP}}] \exists i \in S \ w'_j = w_j^{(i)} \right\}.$$

Informally, for every $w' \in F(W_S)$, every bit w'_j appears somewhere in the j -th column of W_S when viewed as a matrix with a codeword in each row.

We define the security of a fingerprinting code $(\text{Gen}_{\text{FP}}, \text{Trace}_{\text{FP}})$ as follows:

Definition 3.3 (Secure Fingerprinting Code). Let $\varepsilon_{\text{FP}}: \mathbb{N} \rightarrow [0, 1]$ and $\ell_{\text{FP}}: \mathbb{N} \rightarrow \mathbb{N}$ be functions. A pair of algorithms $(\text{Gen}_{\text{FP}}, \text{Trace}_{\text{FP}})$ is a $(\varepsilon_{\text{FP}}, \ell_{\text{FP}})$ -*fingerprinting code* if $\text{Gen}_{\text{FP}}(1^n)$ outputs a codebook $W \in \{0, 1\}^{n \times \ell_{\text{FP}}(n)}$, and furthermore, for every (possibly inefficient) algorithm \mathcal{A}_{FP} , and every non-empty $S \subseteq [n]$,

$$\Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n) \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}(W_S)}} [w' \in F(W_S) \wedge \text{Trace}_{\text{FP}}(W, w') \notin S] \leq \varepsilon_{\text{FP}}(n)$$

Tardos [Tar08] gave a construction of fingerprinting codes of essentially optimal length, improving on the original construction of Boneh and Shaw [BS98].

Theorem 3.4 ([Tar08]). *For every function $\varepsilon_{\text{FP}}: \mathbb{N} \rightarrow [0, 1]$, there exists an $(\varepsilon_{\text{FP}}, 100n^2 \log(n/\varepsilon_{\text{FP}}))$ -fingerprinting code. In particular, there exists a $(1/n^3, 400n^2 \log n)$ -fingerprinting code.*

3.3 Encryption Schemes

We will build our traitor-tracing scheme from a suitable encryption scheme. An encryption scheme is tuple of efficient algorithms $(\text{Gen}, \text{Enc}, \text{Dec})$. All the algorithms may be randomized except for Dec . The scheme has the following syntactic properties:

- The algorithm Gen takes a security parameter κ , runs in time $\text{poly}(\kappa)$, and returns a private key $sk \in \{0, 1\}^\kappa$. Formally $sk \leftarrow_{\text{R}} \text{Gen}(1^\kappa)$.
- The algorithm Enc takes a private key and a message bit $b \in \{0, 1\}$ and generates a ciphertext $c \in \mathcal{C} = \mathcal{C}^{(\kappa)}$. Formally, $c \leftarrow_{\text{R}} \text{Enc}(sk, b)$.
- The algorithm Dec takes a private key $sk \in \{0, 1\}^\kappa$ and a ciphertext $c \in \mathcal{C}^{(\kappa)}$, runs in time $\text{poly}(\kappa)$, and returns a message bit \bar{b} .

First we define (perfectly) correct decryption⁴

Definition 3.5 (Correctness). An encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ is *(perfectly) correct* if for every $b \in \{0, 1\}$, and every $\kappa \in \mathbb{N}$,

$$\Pr_{\substack{sk \leftarrow_{\text{R}} \text{Gen}(1^\kappa) \\ c \leftarrow_{\text{R}} \text{Enc}(sk, b)}} [\text{Dec}(sk, c) = b] = 1.$$

We require that our schemes have the following k_{Enc} -message security property.

Definition 3.6 (Security of Encryption). Let $\varepsilon_{\text{Enc}}: \mathbb{N} \rightarrow [0, 1]$ and $k_{\text{Enc}}, T_{\text{Enc}}: \mathbb{N} \rightarrow \mathbb{N}$ be functions. An encryption scheme $\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})$ is an $(\varepsilon_{\text{Enc}}, k_{\text{Enc}}, T_{\text{Enc}})$ -secure encryption scheme if for every $T_{\text{Enc}}(\kappa, k_{\text{Enc}}(\kappa))$ -time algorithm \mathcal{A}_{Enc} and every $b = (b_1, \dots, b_{k_{\text{Enc}}}), b' = (b'_1, \dots, b'_{k_{\text{Enc}}}) \in \{0, 1\}^{k_{\text{Enc}}}$ (for $k_{\text{Enc}} = k_{\text{Enc}}(\kappa)$),

$$\left| \Pr_{\substack{sk \leftarrow_{\text{R}} \text{Gen}(1^\kappa) \\ c_1, \dots, c_{k_{\text{Enc}}} \leftarrow_{\text{R}} \text{Enc}(sk, b_1), \dots, \text{Enc}(sk, b_{k_{\text{Enc}}})}} [\mathcal{A}_{\text{Enc}}(c_1, \dots, c_{k_{\text{Enc}}}) = 1] - \Pr_{\substack{sk \leftarrow_{\text{R}} \text{Gen}(1^\kappa) \\ c'_1, \dots, c'_{k_{\text{Enc}}} \leftarrow_{\text{R}} \text{Enc}(sk, b'_1), \dots, \text{Enc}(sk, b'_{k_{\text{Enc}}})}} [\mathcal{A}_{\text{Enc}}(c'_1, \dots, c'_{k_{\text{Enc}}}) = 1] \right| \leq \varepsilon_{\text{Enc}}(\kappa)$$

where $c_1, \dots, c_{k_{\text{Enc}}}$ are chosen randomly as $\text{Enc}(sk, b_1), \dots, \text{Enc}(sk, b_{k_{\text{Enc}}})$ and $c'_1, \dots, c'_{k_{\text{Enc}}}$ are chosen randomly as $\text{Enc}(sk, b'_1), \dots, \text{Enc}(sk, b'_{k_{\text{Enc}}})$

Notice that we do not require Π_{Enc} to be secure against adversaries that are given $\text{Enc}(sk, \cdot)$ as an oracle. That is, we do not require CPA security.

Definition 3.7 (Weak Encryption Scheme). A tuple of algorithms $\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})$ is an $(\varepsilon_{\text{Enc}}, k_{\text{Enc}}, T_{\text{Enc}})$ -encryption scheme if it satisfies correctness and $(\varepsilon_{\text{Enc}}, k_{\text{Enc}}, T_{\text{Enc}})$ -security.

⁴It would not substantially affect our results if Dec were allowed to fail with negligible probability, however we will assume perfect correctness for ease of presentation.

3.4 Decryption Function Families

For Theorem 1.2, we want to consider encryption and traitor-tracing schemes where Dec is a “simple” function of the user key (for every ciphertext $c \in \mathcal{C}$ and user key $sk \in \{0, 1\}^\kappa$).

Definition 3.8 (Decryption Function Family). Let $(\text{Gen}, \text{Enc}, \text{Dec})$ be a tuple of algorithms where Gen produces keys in $\{0, 1\}^\kappa$ and Enc produce ciphertexts in $\mathcal{C} = \mathcal{C}^{(\kappa)}$. For every $c \in \mathcal{C}$, we define the c -decryption function $f_c: \{0, 1\}^\kappa \rightarrow \{0, 1\}$ to be $f_c(sk) = \text{Dec}(sk, c)$. We define the *decryption function family* $\mathcal{F}_{\text{Dec}, \kappa} = \{f_c\}_{c \in \mathcal{C}^{(\kappa)}}$.

For a traitor-tracing scheme $(\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$, the family $\mathcal{F}_{\text{Dec}_{\text{TT}}, \kappa}$ is defined analogously.

In what follows, we will say that Π_{Enc} (resp. Π_{TT}) is an encryption scheme (resp. traitor-tracing scheme) with decryption function family $\mathcal{F}_{\text{Dec}, \kappa}$.

4 Attacking Efficient Generic Sanitizers

In this section we will prove our main result, showing that the existence of traitor-tracing schemes (in the sense of Definition 3.2) implies that efficient generic sanitizers cannot guarantee differential privacy.

Theorem 4.1 (Attacking Efficient Generic Sanitizers). *Assume there exists an (n, k_{TT}) -secure traitor-tracing scheme $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ with the decryption function family $\mathcal{F}_{\text{Dec}_{\text{TT}}, \kappa}$. Let $\mathcal{Q} = \bigcup_{d \in \mathbb{N}} \mathcal{Q}^{(d)}$ be any query set such that $\mathcal{F}_{\text{Dec}, d} \subseteq \mathcal{Q}^{(d)}$ for every $d \in \mathbb{N}$. Then there does not exist any sanitizer \mathcal{M} that is simultaneously 1) differentially private, 2) efficient, and 3) $(\mathcal{Q}, k_{\text{TT}}(d))$ -accurate.*

In the typical setting of parameters, $n(\kappa) = \text{poly}(\kappa)$, $k_{\text{TT}}(\kappa) = \tilde{\Theta}(n^2)$, and decryption can be implemented by circuits of size $\text{poly}(n) = \text{poly}(d)$. Then Theorem 4.1 will state that there is no sanitizer \mathcal{M} that takes a database $D \in (\{0, 1\}^d)^{\text{poly}(d)}$, runs in $\text{poly}(d)$ time, and answers $\tilde{\Theta}(n^2)$ queries implemented by circuits of size $\text{poly}(d)$, while satisfying differential privacy.

We now give a sketch of the proof, which follows the approach of Dwork et al. [DNR⁺09]: For every ciphertext $c \in \mathcal{C}^{(d)}$, consider the query $q_c(x) = \text{Dec}(x, c)$ and let $\mathcal{Q}^{(d)}$ contain all of these queries. Notice that $\mathcal{Q}^{(\kappa)} = \mathcal{F}_{\text{Dec}_{\text{TT}}, d}$, and assume there is an efficient mechanism is that is $(\mathcal{Q}^{(d)}, k_{\text{TT}}(d))$ -accurate for these queries. The fact that \mathcal{M} is accurate for these queries will imply that (after small modifications) \mathcal{M} is a k_{TT} -useful pirate decoder (Definition 3.1).

Now consider two experiments: In the first, we construct an n -row database D by running $\text{Gen}_{\text{TT}}(1^d)$ to obtain n user keys, and putting one in each row of D . Then we run Trace_{TT} on $\mathcal{M}(D, \cdot)$ and obtain a user i . Since \mathcal{M} is useful, and Π_{TT} is secure, we will have that $i \in [n]$ with probability close to 1, and thus there is an $i^* \in [n]$ such that $i = i^*$ with probability close to $1/n$.

In the second experiment, we construct a database D' exactly as in the first, however we exclude the key $sk^{(i^*)}$. Since D and D' differ in only one row, differential privacy requires that Trace_{TT} , run with oracle $\mathcal{M}(D', \cdot)$, still outputs i^* with probability close to $1/n$. However, in this experiment, $i^*, sk^{(i^*)}$ is no longer given to the pirate decoder, and thus security of Π_{TT} says that Trace_{TT} , run with this oracle, must output i^* with probability $o(1/n)$. Thus, we will obtain a contradiction.

Proof. Let $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ be the assumed traitor-tracing scheme. For every $d \in \mathbb{N}$, we define the query set

$$\mathcal{Q}_{\text{Dec}_{\text{TT}}}^{(d)} = \left\{ q_c(x) = \text{Dec}_{\text{TT}}(x, c) \mid c \in \mathcal{C}^{(d)} \right\} = \mathcal{F}_{\text{Dec}_{\text{TT}}, d}$$

over $\{0, 1\}^d$ and let $\mathcal{Q}_{\text{Dec}_{\text{TT}}} = \left\{ \mathcal{Q}_{\text{Dec}_{\text{TT}}}^{(d)} \right\}_{d \in \mathbb{N}}$. Recall that since Dec_{TT} runs in time $\text{poly}(d, n(d))$, (at a minimum) $\mathcal{Q}_{\text{Dec}_{\text{TT}}}^{(d)} \subseteq \mathcal{Q}_s^{(d)}$ for some $s(d) = \text{poly}(d, n(d))$.

Assume there exists an efficient, differentially private, $(\mathcal{Q}_{\text{Dec}_{\text{TT}}}, k_{\text{TT}})$ -accurate sanitizer \mathcal{M} . We define the pirate decoder $\mathcal{P}_{\mathcal{M}}$ as follows:

Algorithm 1 The pirate decoder $\mathcal{P}_{\mathcal{M}}$

Input: A set of user keys $(\vec{sk}_S) \in \{0, 1\}^d$ and a set of ciphertexts $c_1, \dots, c_{k_{\text{TT}}}$ ($k_{\text{TT}} = k_{\text{TT}}(d)$).

Construct circuits specifying the queries $q_{c_1}, \dots, q_{c_{k_{\text{TT}}}} \in \mathcal{Q}_{\text{Dec}_{\text{TT}}}^{(d)}$.

Construct a database $D = (sk^{(i)})_{i \in S} \in (\{0, 1\}^d)^{|S|}$.

Let $a_1, \dots, a_{k_{\text{TT}}} \leftarrow_{\mathcal{R}} \mathcal{M}(D, q_{c_1}, \dots, q_{c_{k_{\text{TT}}}})$.

Round the answers $a_1, \dots, a_{k_{\text{TT}}}$ to $\{0, 1\}$ to obtain $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}} \in \{0, 1\}$.

Output: $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}}$.

Recall that if \mathcal{M} is efficient, it runs in time $\text{poly}(d, n(d), k_{\text{TT}}(d), s(d)) = \text{poly}(d, n(d), k_{\text{TT}}(d))$. In this case $\mathcal{P}_{\mathcal{M}}$ runs in time $\text{poly}(d, n(d), k_{\text{TT}}(d))$ as well. To see this, observe that since Dec_{TT} runs in time $s(d) = \text{poly}(d, n(d))$, we can construct a circuit implementing $\text{Dec}_{\text{TT}}(sk, c) = q_c(sk)$ in time $\text{poly}(d, n(d))$. Additionally, the final rounding step can be performed in time $\text{poly}(k_{\text{TT}}(d))$.

Next, we claim that if \mathcal{M} is accurate for $\mathcal{Q}_{\text{Dec}_{\text{TT}}}$, then $\mathcal{P}_{\mathcal{M}}$ is a useful pirate decoder.

Claim 4.2. *If \mathcal{M} is $(\mathcal{Q}, k_{\text{TT}})$ -accurate for some $\mathcal{Q} = \bigcup_{d \in \mathbb{N}} \mathcal{Q}^{(d)}$ such that $\mathcal{Q}_{\text{Dec}_{\text{TT}}}^{(d)} \subseteq \mathcal{Q}^{(d)}$ for every $d \in \mathbb{N}$, then $\mathcal{P}_{\mathcal{M}}$ is a k_{TT} -useful pirate decoder.*

Proof of Claim 4.2. Let $\vec{sk} \in \{0, 1\}^d$ be a set of user keys for Π_{TT} and let $S \subseteq [n]$ be a subset of the users such that $|S| \geq n - 1$. Suppose c and b are such that for every $i \in S$, $\text{Dec}_{\text{TT}}(sk^{(i)}, c) = b$. Then we have that, for D as in $\mathcal{P}_{\mathcal{M}}$,

$$q_c(D) = \frac{1}{|S|} \sum_{i \in S} q_c(sk^{(i)}) = \frac{1}{|S|} \sum_{i \in S} \text{Dec}_{\text{TT}}(sk^{(i)}, c) = b$$

Let $c_1, \dots, c_{k_{\text{TT}}}$ be a set of ciphertexts and let $q_{c_1}, \dots, q_{c_{k_{\text{TT}}}}$ be as in $\mathcal{P}_{\mathcal{M}}$. Let $a_1, \dots, a_{k_{\text{TT}}} \leftarrow_{\mathcal{R}} \mathcal{M}(D, q_{c_1}, \dots, q_{c_{k_{\text{TT}}}})$. The accuracy of \mathcal{M} (with constant error $\alpha < 1/2$) guarantees that

$$\begin{aligned} & \Pr [\exists j \in [k_{\text{TT}}], |a_j - q_{c_j}(D)| \geq 1/2] = o(1/|S|^2) \\ \implies & o(1/n^2) = \Pr [\exists j \in [k_{\text{TT}}], |a_j - q_{c_j}(D)| \geq 1/2] \\ & = \Pr [\exists j \in [k_{\text{TT}}], b \in \{0, 1\}, ((q_{c_j}(D) = b) \wedge (|a_j - b| \geq 1/2))] \\ & = \Pr [\exists j \in [k_{\text{TT}}], b \in \{0, 1\}, ((q_{c_j}(D) = b) \wedge (\text{round}(a_j) \neq b))] \\ & = \Pr \left[\begin{array}{c} \exists j \in [k_{\text{TT}}], b \in \{0, 1\} \\ \left((\forall i \in S, \text{Dec}_{\text{TT}}(sk^{(i)}, c_j) = b) \wedge \left(\mathcal{P}_{\mathcal{M}}(\vec{sk}_S, c_1, \dots, c_{k_{\text{TT}}})_j \neq b \right) \right) \end{array} \right] \end{aligned}$$

Thus, $\mathcal{P}_{\mathcal{M}}$ is k_{TT} -useful. This completes the proof of the claim. \square

Since $\mathcal{P}_{\mathcal{M}}$ is a k_{TT} -useful pirate decoder, and Π_{TT} is a (n, k_{TT}) -secure traitor-tracing scheme, we have that

$$\Pr_{\substack{\vec{s}k \leftarrow \text{RGen}_{\text{TT}}(1^\kappa) \\ \mathcal{P}_{\mathcal{M}}\text{'s, Trace}_{\text{TT}}\text{'s coins}}} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}_{\mathcal{M}}(\vec{s}k, \cdot)}(\vec{s}k) \in [n(\kappa)] \right] \geq 1 - o\left(\frac{1}{n(\kappa)}\right)$$

Thus, for every $\kappa \in \mathbb{N}$, there exists $i^*(\kappa) \in [n(\kappa)]$ such that,

$$\Pr_{\substack{\vec{s}k \leftarrow \text{RGen}_{\text{TT}}(1^\kappa) \\ \mathcal{P}_{\mathcal{M}}\text{'s, Trace}_{\text{TT}}\text{'s coins}}} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}_{\mathcal{M}}(\vec{s}k, \cdot)}(\vec{s}k) = i^*(\kappa) \right] \geq \frac{1}{n(\kappa)} - o\left(\frac{1}{n(\kappa)}\right). \quad (1)$$

Let $S(\kappa) = [n(\kappa)] \setminus \{i^*(\kappa)\}$. Now we claim that if \mathcal{M} is differentially private, then Trace_{TT} will output $i^*(\kappa)$ with significant probability, even $\mathcal{P}_{\mathcal{M}}$ is given the set of keys $\vec{s}k_{S(\kappa)}$, rather than $\vec{s}k$.

Claim 4.3. *If \mathcal{M} is differentially private, then*

$$\Pr_{\substack{\vec{s}k \leftarrow \text{RGen}_{\text{TT}}(1^\kappa) \\ \mathcal{P}_{\mathcal{M}}\text{'s, Trace}_{\text{TT}}\text{'s coins}}} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}_{\mathcal{M}}(\vec{s}k_{S(\kappa)}, \cdot)}(\vec{s}k) = i^*(\kappa) \right] \geq \Omega\left(\frac{1}{n(\kappa)}\right).$$

Proof of Claim 4.3. Fix any κ and let $k_{\text{TT}} = k_{\text{TT}}(\kappa)$ and $i^* = i^*(\kappa)$, $S = S(\kappa)$ as above. Also fix any $\vec{s}k, c_1, \dots, c_{k_{\text{TT}}}$, and let $q_{c_1}, \dots, q_{c_{k_{\text{TT}}}}$ be the queries constructed in the execution of $\mathcal{P}_{\mathcal{M}}$ with $c_1, \dots, c_{k_{\text{TT}}}$ as input. Let $D = \vec{s}k$ and $D_{-i^*} = (\vec{s}k_S)$. By differential privacy (for $\varepsilon = O(1)$, $\delta = o(1/n)$), we have that for every $T \subseteq \mathbb{R}^{k_{\text{TT}}}$

$$\Pr \left[\mathcal{M}(D, q_{c_1}, \dots, q_{c_{k_{\text{TT}}}}) \in T \right] \leq e^{O(1)} \cdot \Pr \left[\mathcal{M}(D_{-i^*}, q_{c_1}, \dots, q_{c_{k_{\text{TT}}}}) \in T \right] + o\left(\frac{1}{n}\right).$$

Note that the queries constructed by $\mathcal{P}_{\mathcal{M}}$ depends only on $c_1, \dots, c_{k_{\text{TT}}}$, not on $\vec{s}k_S$. Also note that the final rounding step does not depend on the input at all. Thus, for every $T \subseteq \{0, 1\}^{k_{\text{TT}}}$

$$\Pr \left[\mathcal{P}_{\mathcal{M}}(\vec{s}k, c_1, \dots, c_{k_{\text{TT}}}) \in T \right] \leq e^{O(1)} \cdot \Pr \left[\mathcal{P}_{\mathcal{M}}(\vec{s}k_S, c_1, \dots, c_{k_{\text{TT}}}) \in T \right] + o\left(\frac{1}{n}\right). \quad (2)$$

Now take $T = T(\vec{s}k, c_1, \dots, c_{k_{\text{TT}}})$ to be the set of responses $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}}$ such that $\text{Trace}_{\text{TT}}(\vec{s}k)$, after querying its oracle on ciphertexts $c_1, \dots, c_{k_{\text{TT}}}$ and receiving responses $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}}$, outputs i^* . The claim now follows by applying (2) to T , averaging over the randomness of $\vec{s}k, c_1, \dots, c_{k_{\text{TT}}}$, and finally combining with (1). \square

To complete the proof, notice that the probability in Claim 4.3 is exactly the probability that Trace_{TT} outputs the user i^* , when given the oracle $\mathcal{P}_{\mathcal{M}}(\vec{s}k_S)$, for $S = [n] \setminus \{i^*\}$. However, the fact that $\mathcal{P}_{\mathcal{M}}$ is efficient, and Π_{TT} is a secure traitor-tracing scheme implies that this probability is $o(1/n)$. Thus we have obtained a contradiction. This completes the proof of the Theorem. \square

5 Constructions of Traitor-Tracing Schemes

In this section we show how to construct traitor-tracing schemes that satisfy Definition 3.2, and thus can be used to instantiate Theorem 4.1. Our construction will make use of an encryption scheme (Definition 3.6) and a fingerprinting code (Definition 3.3). First we will state the construction and

then prove its security (Section 5.1), using the encryption scheme only as a black box. Then we will consider some possible choices for the encryption scheme, and analyze the resulting decryption function family for the traitor-tracing scheme (Section 5.2).

The encryption and decryption functionality of our traitor-tracing scheme, Π_{TT} , is a standard construction: Let Π_{Enc} be an encryption scheme. To generate keys for Π_{TT} , generate n independent user keys $sk^{(1)}, \dots, sk^{(n)}$ for Π_{Enc} , and give one to each user. To encrypt a bit b to all of the users, encrypt it using Π_{Enc} under each user key $sk^{(i)}$ and concatenate the ciphertexts. To decrypt a ciphertext using user key $sk^{(i)}$, find the i -th block of the ciphertext and decrypt it using Π_{Enc} and user key $sk^{(i)}$. We give a more formal specification of the construction below. It will be clear from the construction that $(\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}})$ satisfies the expected properties of encryption.

Algorithm 2 The algorithms $(\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}})$ for Π_{TT} . Let an encryption $\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})$ and a function $n: \mathbb{N} \rightarrow \mathbb{N}$ be parameters of the scheme.

$\text{Gen}_{\text{TT}}(1^\kappa)$:
Let: $n = n(\kappa)$, $\bar{\kappa} = \kappa - \lceil \log n \rceil$
For: $i = 1, 2, \dots, n$
 $\overline{sk}^{(i)} \leftarrow_{\text{R}} \text{Gen}(1^{\bar{\kappa}})$
Let: $sk^{(i)} = (i, \overline{sk}^{(i)})$
EndFor
Output: $(sk^{(1)}, \dots, sk^{(n)})$

$\text{Enc}_{\text{TT}}(sk^{(1)}, \dots, sk^{(n)}, b)$:
For: $i = 1, 2, \dots, n$
Let: $(i, \overline{sk}^{(i)}) = sk^{(i)}$
 $c^{(i)} \leftarrow_{\text{R}} \text{Enc}(\overline{sk}^{(i)}, b)$
EndFor
Output: $c = (c^{(1)}, \dots, c^{(n)})$

$\text{Dec}_{\text{TT}}(sk, c)$:
Let: $c = (c^{(1)}, \dots, c^{(n)})$
Let: $(i, \overline{sk}^{(i)}) = sk$
Output: $\bar{b} = \text{Dec}(\overline{sk}^{(i)}, c^{(i)})$

Although the construction of the scheme is standard, we will need to show that it can be traced in our model using only a small number of non-adaptive queries. First we define a subroutine, TrEnc_{TT} , that “encrypts a matrix” $B \in \{0, 1\}^{n \times k}$. We will use $\vec{b}_j \in \{0, 1\}^n$ to denote the j -th column of B , $\vec{b}^{(i)} \in \{0, 1\}^k$ to denote the i -th row of B , and $b_j^{(i)}$ to denote the (i, j) -th bit of B .

The algorithm will generate k ciphertexts c_1, \dots, c_k . Each ciphertext c_j corresponds to a column of B , and each block of the ciphertext c_j contains an encryption using Enc and $sk^{(i)}$ of the bit $b_j^{(i)}$. Notice that if $k = 1$ (the matrix has only one column) and every row contains the same bit b , then $\text{TrEnc}_{\text{TT}}(\vec{sk}, B)$ is distributed identically to $\text{Enc}(\vec{sk}, b)$.

Now we can specify the tracing algorithm for Π_{TT} . Let Π_{FP} be a fingerprinting code. First, Trace_{TT} will generate a codebook W for the fingerprinting code and then run $\text{TrEnc}_{\text{TT}}(\vec{sk}, W)$ to obtain its challenge ciphertexts. The, Trace_{TT} queries its oracle on these ciphertexts and receives

Algorithm 3 The subroutine TrEnc_{TT} .

$\text{TrEnc}_{\text{TT}}(sk^{(1)}, \dots, sk^{(n)}, B)$:
 Let: $B = (\vec{b}_1, \dots, \vec{b}_k)$ be the columns of B .
 For: $j = 1, \dots, k$
 For: $i = 1, \dots, n$
 Let: $(i, \vec{sk}^{(i)}) = sk^{(i)}$
 $c_j^{(i)} \leftarrow \text{Enc}(\vec{sk}_i, \vec{b}_j^{(i)})$
 EndFor
 Let: $c_j = (c_j^{(1)}, \dots, c_j^{(n)})$
 EndFor
Output: c_1, \dots, c_k

$\vec{b}_1, \dots, \vec{b}_{\ell_{\text{FP}}}$ in response. Finally, it treats these responses as a word w' and runs the tracing algorithm for the fingerprinting code, Trace_{FP} on w' , and repeats the output of Trace_{FP} as its own output.

Algorithm 4 The tracing algorithm for Π_{TT} . Let a length $\ell_{\text{FP}} = \ell_{\text{FP}}(n)$ fingerprinting code $\Pi_{\text{FP}} = (\text{Gen}_{\text{FP}}, \text{Trace}_{\text{FP}})$ and an encryption scheme $\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})$ be parameters of the scheme.

$\text{Trace}_{\text{TT}}^{\mathcal{P}}(\vec{sk})$:
 Let: $n = n(\kappa)$, $\ell_{\text{FP}} = \ell_{\text{FP}}(n)$
 Let: $W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n)$
 Let: $c_1, \dots, c_{\ell_{\text{FP}}} \leftarrow_{\text{R}} \text{TrEnc}_{\text{TT}}(\vec{sk}, W)$
 Let: $\vec{b}_1, \dots, \vec{b}_{\ell_{\text{FP}}} \leftarrow_{\text{R}} \mathcal{P}(c_1, \dots, c_{\ell_{\text{FP}}})$
 Let: $i \leftarrow_{\text{R}} \text{Trace}_{\text{FP}}(W, w')$ where $w' = \vec{b}_1 \| \dots \| \vec{b}_{\ell_{\text{FP}}}$ is the concatenation of $\vec{b}_1, \dots, \vec{b}_{\ell_{\text{FP}}}$
 Output: i

5.1 Security of Π_{TT}

In this section we will prove that our construction of $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ is an $(n, \ell_{\text{FP}}(n))$ -secure traitor-tracing scheme. It can be verified from the specification of the scheme that it has the desired syntactic properties, that it generates $n(\kappa)$ user keys, and that the tracing algorithm makes $\ell_{\text{FP}}(n(\kappa))$ non-adaptive queries to its oracle.

Before proving the security of the scheme, we will give some intuition for why a pirate decoder $\mathcal{P}(\vec{sk}_S, \cdot)$ can be successfully traced. The algorithm Trace_{TT} is going to generate queries $c_1, \dots, c_{\ell_{\text{FP}}}$ using $\text{TrEnc}_{\text{TT}}(\vec{sk}, W)$, for a fingerprinting codebook W . By the construction of TrEnc_{TT} and the correctness of Π_{Enc} , if $j \in \text{Crit}(W_S)$, then every user $i \in S$ will decrypt c_j to the same bit $b_j = w_j^{(i)}$. If \mathcal{P} is a ℓ_{FP} -useful pirate decoder, then with high probability it answers the queries with $\vec{b}_1, \dots, \vec{b}_{\ell_{\text{FP}}}$ such that for every $j \in \text{Crit}(W_S)$, $\vec{b}_j = b_j$. Thus w' will be in $F(W_S)$ with high probability. The security of the fingerprinting code now implies that $i \in S$ with high probability.

The catch in this argument is that TrEnc_{TT} takes all of W as input, however an attacker for the fingerprinting code is only allowed to see W_S , and thus cannot simulate TrEnc_{TT} in a security reduction. However, if \mathcal{P} only has keys \vec{sk}_S , and $i \notin S$, then an efficient \mathcal{P} cannot decrypt the i -th block of a ciphertext $c = (c^{(1)}, \dots, c^{(n)})$. Notice that the ciphertext components $c_1^{(i)}, \dots, c_{\ell_{\text{FP}}}^{(i)}$ are

exactly those that contain encryptions of the i -th row of W , which is the i -th codeword in W . Thus the i -th codeword is computationally hidden from \mathcal{P} anyway, and we could replace that codeword with the all-zeros codeword without significantly affecting the success probability of \mathcal{P} . Formalizing this intuition will yield a valid attacker for the fingerprinting code, and obtain a contradiction.

We now formalize this intuition and prove the following theorem:

Theorem 5.1 (From Encryption to Traitor-Tracing). *Assume that there exists an $(\varepsilon_{\text{Enc}}, k_{\text{Enc}}, T_{\text{Enc}})$ -secure encryption scheme, Π_{Enc} and an $(\varepsilon_{\text{FP}}, \ell_{\text{FP}})$ -fingerprinting code, Π_{FP} . Let $n: \mathbb{N} \rightarrow \mathbb{N}$ and $k_{\text{TT}}: \mathbb{N} \rightarrow [0, 1]$ be any functions such that for every $\kappa \in \mathbb{N}$*

1. $n(\kappa) \cdot \varepsilon_{\text{Enc}}(\kappa) + \varepsilon_{\text{FP}}(n(\kappa)) = o\left(\frac{1}{n(\kappa)^2}\right)$,
2. $k_{\text{Enc}}(\kappa) \geq \ell_{\text{FP}}(n(\kappa)) = k_{\text{TT}}(\kappa)$, and
3. $\text{poly}(\kappa, n(\kappa), k_{\text{TT}}(\kappa)) \leq T_{\text{Enc}}(\kappa - \lceil \log(n(\kappa)) \rceil, k_{\text{TT}}(\kappa))$

Then $\Pi_{\text{TT}} = (\text{Gen}_{\text{TT}}, \text{Enc}_{\text{TT}}, \text{Dec}_{\text{TT}}, \text{Trace}_{\text{TT}})$ with parameters $n, k_{\text{TT}}, \Pi_{\text{Enc}}, \Pi_{\text{FP}}$ is a (n, k_{TT}) -traitor-tracing scheme.

Proof. Suppose there exists a $\text{poly}(\kappa, n(\kappa), k_{\text{TT}}(\kappa))$ -time pirate decoder \mathcal{P} that violates the security of Π_{TT} . There is, for every $\kappa \in \mathbb{N}$, there exists $S = S(\kappa) \subseteq [n(\kappa)]$, $|S(\kappa)| \geq n(\kappa) - 1$, such that

$$\Pr_{\vec{s}k \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa)} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}(\vec{s}k_{S(\kappa)}, \cdot)}(\vec{s}k) \notin S \right] = \Omega\left(\frac{1}{n(\kappa)}\right)$$

where the probability is also taken over the coins of \mathcal{P} and Trace_{TT} . Thus, for a randomly chosen $S(\kappa) \subseteq [n(\kappa)]$ such that $|S(\kappa)| \geq n(\kappa) - 1$,

$$\Pr_{\vec{s}k \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa), S(\kappa)} \left[\text{Trace}_{\text{TT}}^{\mathcal{P}(\vec{s}k_{S(\kappa)}, \cdot)}(\vec{s}k) \notin S \right] = \Omega\left(\frac{1}{n(\kappa)^2}\right)$$

where this probability is also taken over the coins of \mathcal{P} and Trace_{TT} . We will show that such a pirate decoder must either violate the security of the encryption scheme or violate the security of the fingerprinting code.

First we define some notation for specifying the adversaries used in the argument. Let $W = (w^{(1)}, \dots, w^{(n)}) \in \{0, 1\}^{n \times \ell_{\text{FP}}}$ be a codebook for the fingerprinting code, represented as a matrix with the i -th row containing the i -th codeword. For $S \subseteq [n]$, $|S| = n - 1$, where $S = [n] \setminus \{i\}$, and codebook W , we write $W_S = (w^{(1)}, \dots, w^{(i-1)}, w^{(i+1)}, \dots, w^{(n)}) \in \{0, 1\}^{(n-1) \times \ell_{\text{FP}}}$ to be the codebook W with the i -th row removed. We also write $\widetilde{W}_S = (w^{(1)}, \dots, w^{(i-1)}, \vec{0}, w^{(i+1)}, \dots, w^{(n)}) \in \{0, 1\}^{n \times \ell_{\text{FP}}}$ to be the codebook W with the i -th row present but replaced with the all-zeros codeword. Note that \widetilde{W}_S can be computed only from W_S , even though W cannot. If $S = [n]$ we will use the same notation, however in this case notice that $W = W_S = \widetilde{W}_S$.

Consider the following algorithm $\mathcal{A}_{\text{FP}}^{\mathcal{P}}$

Since $\mathcal{A}_{\text{FP}}^{\mathcal{P}}$ is a valid attacker in the fingerprinting security game for every \mathcal{P} (it only takes (S, W_S) as input), for every $S \subseteq [n]$, $S \neq \emptyset$,

$$\Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n) \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)}} \left[w' \in F(W_S) \wedge \text{Trace}_{\text{FP}}(W, w') \notin S \right] \leq \varepsilon_{\text{FP}}(n) = \varepsilon_{\text{FP}}(n(\kappa))$$

Algorithm 5 The fingerprinting security adversary $\mathcal{A}_{\text{FP}}^{\mathcal{P}}$.

$\mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)$:

Let n be the number of users for the fingerprinting code, and choose κ such that $n(\kappa) = n$

Let: $\vec{sk} = (sk^{(1)}, \dots, sk^{(n)}) \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa)$

Let: $(c_1, \dots, c_{\ell_{\text{FP}}}) = \text{TrEnc}_{\text{TT}}(\vec{sk}, \widetilde{W}_S)$

Let: $(\bar{b}_1, \dots, \bar{b}_{\ell_{\text{FP}}}) = \mathcal{P}(\vec{sk}_S, c_1, \dots, c_{\ell_{\text{FP}}})$.

Output: $w' = \bar{b}_1 \| \dots \| \bar{b}_{\ell_{\text{FP}}}$, the concatenation of $\bar{b}_1, \dots, \bar{b}_{\ell_{\text{FP}}}$.

Thus, for randomly chosen $S = S(\kappa) \subseteq [n(\kappa)]$, $|S| \geq n(\kappa) - 1$,

$$\Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)}} [w' \in F(W_S) \wedge \text{Trace}_{\text{FP}}(W, w') \notin S] \leq \varepsilon_{\text{FP}}(n) = \varepsilon_{\text{FP}}(n(\kappa)) \quad (3)$$

The next claim states that if \mathcal{P} is a ℓ_{FP} -useful pirate decoder, then $w' \in F(W_S)$ with high probability.

Claim 5.2. *Let $k_{\text{TT}} = k_{\text{TT}}(\kappa) = \ell_{\text{FP}}(n(\kappa))$ for every $\kappa \in \mathbb{N}$. If \mathcal{P} is a k_{TT} -useful pirate decoder, then for every $\kappa \in \mathbb{N}$, $S = S(\kappa) \subseteq [n(\kappa)]$ such that $|S| \geq n(\kappa) - 1$, and $W \in \{0, 1\}^{n \times \ell_{\text{FP}}(n)}$ (for $n = n(\kappa)$)*

$$\Pr_{w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)} [w' \notin F(W_S)] = o\left(\frac{1}{n(\kappa)^2}\right)$$

Proof of Claim 5.2. If \mathcal{P} is k_{TT} -useful, then, in particular, for any $\vec{sk} = (sk^{(1)}, \dots, sk^{(n)})$, any $S \subseteq [n]$ such that $|S| \geq n - 1$, any ciphertexts $c_1, \dots, c_{k_{\text{TT}}}$, and $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}} \leftarrow_{\text{R}} \mathcal{P}(\vec{sk}_S, c_1, \dots, c_{k_{\text{TT}}})$.

$$\Pr [\exists j \in [k_{\text{TT}}], b \in \{0, 1\} \left(\left(\forall i \in S, \text{Dec}_{\text{TT}}(sk^{(i)}, c_j) = b \right) \wedge (\bar{b}_j \neq b) \right)] = o\left(\frac{1}{n(\kappa)^2}\right) \quad (4)$$

Consider any $j \in \text{Crit}(W_S)$. There exists b_j such that for every $i \in S$, $w_j^{(i)} = b_j$. Let $c_1, \dots, c_{\ell_{\text{FP}}} \leftarrow_{\text{R}} \text{TrEnc}_{\text{TT}}(\vec{sk}, \widetilde{W}_S)$ and $c_j = (c_j^{(1)}, \dots, c_j^{(n)})$. By construction of TrEnc_{TT} , for every $i \in S$, $c_j^{(i)}$ is distributed as $\text{Enc}(sk^{(i)}, w_j^{(i)}) = \text{Enc}(sk^{(i)}, b_j)$. Thus, by the correctness of Π_{Enc} , for every $i \in S$, $\text{Dec}(sk^{(i)}, c_j^{(i)}) = b_j$, and by the construction of Π_{TT} , for every $i \in S$, $\text{Dec}_{\text{TT}}(sk^{(i)}, c_j) = b_j$. Thus, by (4),

$$\begin{aligned} & \Pr_{w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)} [w' \notin F(W_S)] \\ &= \Pr [\exists j \in \text{Crit}(W_S), i \in S \text{ s.t. } \bar{b}_j \neq b_j] \\ &= \Pr [\exists j \in [k_{\text{TT}}], \left(\left(\forall i \in S, \text{Dec}_{\text{TT}}(sk^{(i)}, c_j) = b_j \right) \wedge (\bar{b}_j \neq b_j) \right)] = o\left(\frac{1}{n(\kappa)^2}\right) \end{aligned}$$

where the latter two probabilities are taken over the random choice of $\vec{sk} \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}(1^\kappa)$, the randomness of the encryptions $c_1, \dots, c_{k_{\text{TT}}} \leftarrow_{\text{R}} \text{TrEnc}_{\text{TT}}(\vec{sk}, \widetilde{W}_S)$, and the randomness of the pirate $\bar{b}_1, \dots, \bar{b}_{k_{\text{TT}}} \leftarrow_{\text{R}} \mathcal{P}(\vec{sk}_S, c_1, \dots, c_{k_{\text{TT}}})$. This completes the proof of the claim. \square

Combining this claim with (3) we obtain

$$\begin{aligned}
\varepsilon_{\text{FP}}(n(\kappa)) &\geq \Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)}} [w' \in F(W_S) \wedge \text{Trace}_{\text{FP}}(W, w') \notin S] \\
&\geq \Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)}} [\text{Trace}_{\text{FP}}(W, w') \notin S] - o\left(\frac{1}{n(\kappa)^2}\right)
\end{aligned} \tag{5}$$

Now consider the algorithm $\tilde{\mathcal{A}}^{\mathcal{P}}(S, W)$, which runs exactly as $\mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)$ but encrypts the matrix W instead of \tilde{W}_S . From the construction of Π_{TT} , it can be verified that for every $S \subseteq [n]$, $|S| \geq n - 1$,

$$\begin{aligned}
\Pr [\text{Trace}_{\text{TT}}^{\mathcal{P}(\vec{sk}_S, \cdot)}(\vec{sk}) \notin S] &\leq n \cdot \Pr_S [\text{Trace}_{\text{TT}}^{\mathcal{P}(\vec{sk}_S, \cdot)}(\vec{sk}) \notin S] \\
&\leq n \cdot \Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \tilde{\mathcal{A}}^{\mathcal{P}}(S, W)}} [\text{Trace}_{\text{FP}}(W, w') \notin S].
\end{aligned}$$

Thus, in order to complete the proof, it will be sufficient to prove the following claim.

Claim 5.3. *If Π_{Enc} is $(\varepsilon_{\text{Enc}}, k_{\text{Enc}}, T_{\text{Enc}})$ -secure for $k_{\text{Enc}}, T_{\text{Enc}}$ as in the statement of the Theorem, then for every $\text{poly}(\kappa, n(\kappa), k_{\text{TT}}(\kappa))$ pirate decoder \mathcal{P} ,*

$$\left| \Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \mathcal{A}_{\text{FP}}^{\mathcal{P}}(S, W_S)}} [\text{Trace}_{\text{FP}}(W, w') \notin S] - \Pr_{\substack{W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n), S \\ w' \leftarrow_{\text{R}} \tilde{\mathcal{A}}^{\mathcal{P}}(S, W)}} [\text{Trace}_{\text{FP}}(W, w') \notin S] \right| \leq \varepsilon_{\text{Enc}}(\kappa)$$

Proof of Claim 5.3. Let $\Pi_{\text{Enc}} = (\text{Gen}, \text{Enc}, \text{Dec})$ be the encryption scheme. Consider a random set of ciphertexts $c_1^{(i)}, \dots, c_{\ell_{\text{FP}}}^{(i)}$ for $\ell_{\text{FP}} = \ell_{\text{FP}}(n(\kappa))$ and for i to be chosen later. The ciphertexts can be generated in one of two ways: In either case, generate a random key $sk^{(i)} \leftarrow_{\text{R}} \text{Gen}(1^\kappa)$.

- In the first case, which we call the “all-zeros case”, $c_1^{(i)} \leftarrow_{\text{R}} \text{Enc}(sk^{(i)}, 0), \dots, c_{k_{\text{TT}}}^{(i)} \leftarrow_{\text{R}} \text{Enc}(sk^{(i)}, 0)$.
- In the second case, which we call the “codeword case”, there is a codeword $w^{(i)} \in \{0, 1\}^{\ell_{\text{FP}}}$ and $c_1^{(i)} \leftarrow_{\text{R}} \text{Enc}(sk^{(i)}, w_1^{(i)}), \dots, c_{\ell_{\text{FP}}}^{(i)} \leftarrow_{\text{R}} \text{Enc}(sk^{(i)}, w_{\ell_{\text{FP}}}^{(i)})$.

Now consider the following algorithm We prove the claim by the following series of observation about \mathcal{A}_{Enc} : 1) The challenge ciphertexts $c_1^{(i)}, \dots, c_{\ell_{\text{FP}}}^{(i)}$ are distributed either as the all-zeros case or the codeword case, as described above. 2) Since Gen_{TT} generates n independent keys from Gen , and in either case of the challenge ciphertexts a key for Gen is chosen independently from Gen , all the keys used in \mathcal{A}_{Enc} have exactly the same distribution as $\vec{sk} \leftarrow_{\text{R}} \text{Gen}_{\text{TT}}$. 3) In the all-zeros case, the ciphertexts $c_1, \dots, c_{\ell_{\text{FP}}}$ are distributed exactly as the ciphertexts in $\mathcal{A}_{\text{FP}}(S, W_S)$, and in the codeword case, the ciphertexts are distributed exactly as the ciphertexts in $\tilde{\mathcal{A}}(S, W)$. 4) If \mathcal{P} runs in $\text{poly}(\kappa, n(\kappa), k_{\text{TT}}(\kappa))$ -time, then so does \mathcal{A}_{Enc} , thus, by the condition on T_{Enc} , \mathcal{A}_{Enc} is a valid attacker for the encryption scheme.

From these observation, we conclude that if the claim is false, then \mathcal{A}_{Enc} violates the security of Π_{Enc} . \square

We now complete the proof of the theorem by combining Equation (5) and Claim 5.3. \square

Algorithm 6 The encryption adversary $\mathcal{A}_{\text{Enc}}^{\mathcal{P}}$

$\mathcal{A}_{\text{Enc}}^{\mathcal{P}}(1^\kappa)$

Let: $n = n(\kappa)$, $\ell_{\text{FP}} = \ell_{\text{FP}}(n)$, $i \leftarrow_{\text{R}} [n]$, let $S = [n] \setminus \{i\}$

Let: choose $n - 1$ userkeys $\vec{sk}_S \leftarrow_{\text{R}} \text{Gen}(1^\kappa)$, independently

Let: $W \leftarrow_{\text{R}} \text{Gen}_{\text{FP}}(1^n)$

Request encryptions for either message $\vec{0} \in \{0, 1\}^{\ell_{\text{FP}}}$ or for $w^{(i)}$ from Π_{Enc}

Receive a sequence of challenge ciphertexts $c_1^{(i)}, \dots, c_{\ell_{\text{FP}}}^{(i)}$

Construct a sequence of ciphertexts $c_1, \dots, c_{\ell_{\text{FP}}}$ mimicking TrEnc_{TT} as follows:

$(c_1^{(-i)}, \dots, c_{\ell_{\text{FP}}}^{(-i)}) \leftarrow_{\text{R}} \text{TrEnc}_{\text{TT}}(\vec{sk}_{-i}, W_{-i})$

For every $j \in [\ell_{\text{FP}}]$, c_j contains $c_j^{(i)}$ in the i -th block and $c_j^{(-i)}$ in all other blocks

Let: $\bar{b}_1, \dots, \bar{b}_{\ell_{\text{FP}}} \leftarrow_{\text{R}} \mathcal{P}(\vec{sk}_S, c_1, \dots, c_{\ell_{\text{FP}}})$

Let: $w' = \bar{b}_1 \parallel \dots \parallel \bar{b}_{\ell_{\text{FP}}}$

Let: $i \leftarrow_{\text{R}} \text{Trace}_{\text{FP}}(W, w')$

Output: 1 if $\text{Trace}_{\text{FP}}(W, w') \notin S$, 0 otherwise.

5.2 Decryption Function Family of Π_{TT}

In this section we consider the complexity of the decryption function used by our traitor-tracing scheme Π_{TT} (Section 5.1). Since Π_{TT} uses an encryption scheme Π_{Enc} as a building block, its complexity will depend on the complexity of the underlying encryption scheme. The following simple lemma gives a description of the decryption function of Π_{TT} as a circuit with gates computing the decryption function for Π_{Enc}

Lemma 5.4 (Decryption Function Family for Π_{TT}). *Let Π_{TT} be as defined, with Π_{Enc} as its underlying encryption scheme. Let $(i, \vec{sk}) = sk \in \{0, 1\}^\kappa$ be any user key for Π_{TT} and let $c = (c^{(1)}, \dots, c^{(n)}) \in \mathcal{C}^{(\kappa)}$ be any ciphertext (for security parameter κ). Then*

$$\text{Dec}_{\text{TT}, c}(sk) = \text{Dec}_{\text{TT}, c}(i, \vec{sk}) = \bigvee_{i' \in [n]} (\mathbf{1}_{i'}(i) \wedge \text{Dec}_{c(i')}(sk))$$

Here, the function $\mathbf{1}_x(y)$ takes the value 1 if $y = x$ and 0 otherwise. The truth of the lemma can easily be verified from the construction of Dec_{TT} . Also note that the function $\mathbf{1}_{i'}: \{0, 1\}^{\lceil \log n \rceil} \rightarrow \{0, 1\}$ is just a conjunction of $\lceil \log n \rceil$ bits, and we need to compute n of these functions. In addition to computing $\mathbf{1}_{i'}$ and $\text{Dec}_{c(i')}$, there are n conjunctions and a single outer disjunction. Thus we add an additional $n + 1$ gates and increase the depth by 2. Hence, an intuitive summary of the lemma is that if Dec can be implemented by circuits of size s and depth h , Dec_{TT} can be implemented by circuits of size $s + \tilde{O}(n)$ and depth $h + 2$. This summary will be precise enough to state our main results.

By combining the preceding Lemma with Theorem 5.1, we can obtain the following corollary.

Corollary 5.5 (One-way Functions Imply traitor-tracing w/ Poly-Time Decryption). *Let $n = n(\kappa)$ be any polynomial in κ . Assuming the existence of one-way functions, there exists an $(n, \tilde{O}(n^2))$ -secure traitor-tracing scheme with decryption $\mathcal{F}_{\text{DecTT}, d} \subseteq \mathcal{Q}_t^{(d)}$ for some function $t = t(d) = \text{poly}(d)$ and every $d \in \mathbb{N}$.*

Proof. The existence of one-way functions implies the existence of an encryption scheme Π_{Enc} that is $(1/\kappa^a, \kappa^a, \kappa^a)$ -secure for every constant $a > 0$ with decryption function $\mathcal{F}_{\text{Dec},d} \subseteq \mathcal{Q}_{t'}^{(d)}$ for some $t' = t'(\kappa) = \text{poly}(\kappa)$ and every $d \in \mathbb{N}$. From Lemma 5.4, it is easy to see that if Π_{TT} uses Π_{Enc} as its encryption scheme, then $\mathcal{F}_{\text{DecTT},d} \subseteq \mathcal{Q}_t^{(d)}$ for $t(d) = t'(d) + \tilde{O}(n(d)) = \text{poly}(d)$. \square

Theorem 1.1 in the introduction follows by combining Theorem 4.1 with Corollary 5.5.

We will now consider the possibility of constructing a traitor-tracing scheme where the decryption functionality can be implemented by circuits of constant depth, and thus obtaining hardness results for generic sanitizers that are efficient for constant-depth queries (Theorem 1.2). We will give a candidate for such a scheme using the notion of local pseudorandom generators.

Definition 5.6 (Local Pseudorandom Generator). An efficient algorithm $G: \{0, 1\}^\kappa \rightarrow \{0, 1\}^{s_{\text{PRG}}(\kappa)}$ is a $(\varepsilon_{\text{PRG}}, s_{\text{PRG}})$ -pseudorandom generator if for every $\text{poly}(s_{\text{PRG}}(\kappa))$ -time adversary \mathcal{A}_{PRG}

$$|\Pr[\mathcal{A}_{\text{PRG}}(G(U_\kappa)) = 1] - \Pr[\mathcal{A}_{\text{PRG}}(U_{s_{\text{PRG}}(\kappa)}) = 1]| \leq \varepsilon_{\text{PRG}}(\kappa)$$

If, in addition, there exists $L \in \mathbb{N}$ such that for every $\kappa \in \mathbb{N}$, $i \in [s_{\text{PRG}}(\kappa)]$, there exists $V_i = \{v_{i,1}, \dots, v_{i,L}\} \subseteq [\kappa]$ and $g_i: \{0, 1\}^L \rightarrow \{0, 1\}$, such that

$$G(s)_i = g_i(s_{v_{i,1}}, s_{v_{i,2}}, \dots, s_{v_{i,L}})$$

then G is a $(\varepsilon_{\text{PRG}}, s_{\text{PRG}}, L)$ -local pseudorandom generator.

It is a well known result in Cryptography that pseudorandom generators imply encryption schemes satisfying Definition 3.6 (for certain ranges of parameters). We will use a particular construction whose decryption can be computed in constant-depth whenever the underlying PRG is locally-computable (or, more generally, computable by constant-depth circuits).

Lemma 5.7 (Local PRGs \rightarrow Encryption). *If there exists a $(\varepsilon_{\text{PRG}}(\kappa), s_{\text{PRG}}(\kappa), L)$ -local pseudorandom generator G , then there exists an $(\varepsilon_{\text{Enc}} = \varepsilon_{\text{PRG}}, k_{\text{Enc}} = \sqrt{s_{\text{PRG}}(\kappa)})$ -Secure Encryption Scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ with $\mathcal{F}_{\text{Dec},d} \subseteq \mathcal{Q}_{t,4}^{(d)}$ for some $t = t(d) = \text{poly}(d)$ and every $d \in \mathbb{N}$.*

The construction is the standard “computational one-time pad”, however we give a construction to verify that the decryption can be computed by constant-depth circuits.

Proof. We construct the scheme as follows. Let $G: \{0, 1\}^\kappa \rightarrow \{0, 1\}^{s_{\text{PRG}}(\kappa)}$ be a pseudorandom generator. Suppose for every $\kappa \in \mathbb{N}$, every bit of G ’s output is computable from L bits of its seed. That is, for every $i \in \{1, 2, \dots, s_{\text{PRG}}(\kappa)\}$ there is a function $g_i: \{0, 1\}^L \rightarrow \{0, 1\}$, and a set $V_i \subseteq [\kappa]$, $|V_i| \leq L$ such that $G(s)_i = g_i(s_{V_i})$. Here s_{V_i} is the restriction of s to the indices in V_i . Let $\mathbf{1}_i(j)$ be the indicator variable for the condition $j = i$. Then for every $c = (r, b) \in \mathcal{C}$, we can write

$$\text{Dec}_{(r,b)}(s) = \bigvee_{i \in [s_{\text{PRG}}(\kappa)]} \mathbf{1}_i(r) \wedge (g_i(s_{V_i}) \oplus b).$$

Observe that, since g_i is a function of ℓ bits of the input, it can be computed by a size- 2^L , depth-2 circuit, thus $g_i(s_{V_i}) \oplus b$ can be computed by a size $2^L + 1$, depth-3 circuit. The indicator $\mathbf{1}_i$ can be computed by a conjunction of $\lceil \log_2 s_{\text{PRG}}(\kappa) \rceil$ bits, which is a size- $\lceil \log_2 s_{\text{PRG}}(\kappa) \rceil$, depth-1 circuit. The outer disjunction increases the depth by one level and the size by 1. Putting it all together, we have shown that $\text{Dec}_{r,b}(s)$ can be computed by depth-4 circuits of size $\tilde{O}(2^L s_{\text{PRG}}(\kappa)) = \text{poly}(s_{\text{PRG}}(\kappa))$. \square

Algorithm 7 An encryption scheme Π_{Enc} that can be decrypted in constant depth.

$\text{Gen}(1^\kappa)$:

$s \leftarrow_{\text{R}} \{0, 1\}^\kappa$

Output: $sk = s$

$\text{Enc}(sk, b)$:

$r \leftarrow_{\text{R}} \{1, 2, \dots, s_{\text{PRG}}(\kappa)\}$

Output: $c = (r, G(sk)_r \oplus b)$

$\text{Dec}(sk, c)$:

$(r', b') = c$

Output: $b = G(sk)_r \oplus b'$

Combining Theorem 5.1 with Lemmas 5.4 and 5.7 yields the following corollary.

Corollary 5.8 (Local Pseudorandom Generators Imply traitor-tracing w/ AC^0 Decryption). *Let $n = n(\kappa)$ be any polynomial in κ . Assuming the existence of a $(o(1/n^2), n^5, L)$ -local pseudorandom generator for some constant $L \in \mathbb{N}$, there exists an $(n, \tilde{O}(n^2))$ -secure traitor-tracing scheme with decryption function family $\mathcal{F}_{\text{DecTT}, d} \subseteq \mathcal{Q}_{t,6}^{(d)}$ for some $t = t(d) = \text{poly}(d)$ and every $d \in \mathbb{N}$.*

Proof. By Lemma 5.7, the assumed pseudorandom generator implies an encryption scheme with $\mathcal{F}_{\text{Dec}, d} \subseteq \mathcal{Q}_{t',4}^{(d)}$ for some $t' = t'(\kappa) = \text{poly}(\kappa)$ and every $d \in \mathbb{N}$. From Lemma 5.4, it is easy to see that if Π_{TT} uses Π_{Enc} as its encryption scheme, then $\mathcal{F}_{\text{DecTT}, d} \subseteq \mathcal{Q}_{t,6}^{(d)}$ for $t(d) = t'(d) + \tilde{O}(n(d)) = \text{poly}(d)$. \square

To support the plausibility of the assumption, we remark that a recent result of Applebaum [App12] gives a candidate construction of a local pseudorandom generators (in the range of parameters we require). We refer the reader to [App12] for more discussion of the plausibility of this assumption.

Theorem 1.2 in the introduction follows by combining Theorem 4.1 with Corollary 5.8.

Acknowledgements

We thank Cynthia Dwork for many inspiring discussions about differential privacy, and in particular for suggesting that we look further at the connection between traitor-tracing and differential privacy. We thank Salil Vadhan for many helpful discussions about the connection between traitor-tracing and differential privacy, and about the presentation of this work. We also thank Dan Boneh, Moritz Hardt, Hart Montgomery, Ananth Raghunathan, Aaron Roth, Guy Rothblum, and Thomas Steinke for helpful discussions.

References

- [App12] Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. In Howard J. Karloff and Toniann Pitassi, editors, *STOC*, pages 805–816. ACM, 2012.

- [BCD⁺07] Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In Leonid Libkin, editor, *PODS*, pages 273–282. ACM, 2007.
- [BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In Chen Li, editor, *PODS*, pages 128–138. ACM, 2005.
- [BLR08] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to non-interactive database privacy. In Cynthia Dwork, editor, *STOC*, pages 609–618. ACM, 2008.
- [BN08] Dan Boneh and Moni Naor. Traitor tracing with constant size ciphertext. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *ACM Conference on Computer and Communications Security*, pages 501–510. ACM, 2008.
- [BS98] Dan Boneh and James Shaw. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory*, 44(5):1897–1905, 1998.
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 573–592. Springer, 2006.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In Yvo Desmedt, editor, *CRYPTO*, volume 839 of *Lecture Notes in Computer Science*, pages 257–270. Springer, 1994.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC ’06*, pages 265–284, 2006.
- [DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210. ACM, 2003.
- [DNR⁺09] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *STOC ’09*, pages 381–390, 2009.
- [DNV12] Cynthia Dwork, Moni Naor, and Salil Vadhan. The privacy of the analyst and the power of the state. *Manuscript*, 2012.
- [DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *FOCS*, pages 51–60. IEEE Computer Society, 2010.
- [GHRU11] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately releasing conjunctions and the statistical query barrier. In *STOC ’11*, pages 803–812, 2011.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In Ronald Cramer, editor, *TCC*, volume 7194 of *Lecture Notes in Computer Science*, pages 339–356. Springer, 2012.
- [HLM10] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for differentially private data release. *CoRR*, abs/1012.4763, 2010.

- [HR10] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS*, pages 61–70. IEEE Computer Society, 2010.
- [HRS12] Moritz Hardt, Guy N. Rothblum, and Rocco A. Servedio. Private data release via learning thresholds. In Dana Randall, editor, *SODA*, pages 168–187. SIAM, 2012.
- [KRSU10] Shiva Prasad Kasiviswanathan, Mark Rudelson, Adam Smith, and Jonathan Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In Schulman [Sch10], pages 775–784.
- [KY01] Aggelos Kiayias and Moti Yung. On crafty pirates and foxy tracers. In Tomas Sander, editor, *Digital Rights Management Workshop*, volume 2320 of *Lecture Notes in Computer Science*, pages 22–39. Springer, 2001.
- [RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *STOC '10*, pages 765–774, 2010.
- [Sch10] Leonard J. Schulman, editor. *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*. ACM, 2010.
- [Tar08] Gábor Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2), 2008.
- [TUV12] Justin Thaler, Jonathan Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *ICALP (1)*, volume 7391 of *Lecture Notes in Computer Science*, pages 810–821. Springer, 2012.
- [UV11] Jonathan Ullman and Salil P. Vadhan. PCPs and the hardness of generating private synthetic data. In *TCC '11*, pages 400–416, 2011.